Methods and Tools of Analysis

1. Introduction

From the moment of their invention, computers have generated complex social, ethical, and value concerns. As computer technology evolves and gets deployed in new ways, certain issues persist – issues of privacy, property rights, accountability, and social values. At the same time, seemingly new and unique issues emerge.

The ethical issues can be organized in at least three different ways:

- According to the type of technology
- According to the sector in which the technology is used.
- According to ethical concepts or themes.

According to the type of technology

This is to organize the ethical issues by type of technology and its use. When computers were first invented, they were understood to be essentially sophisticated calculating machines but they seemed to have the capacity to do that which was thought to be uniquely human -- to reason and exhibit a high degree of rationality; hence, there was concern that computers threatened ideas about what it means to be human. In the shadow of World War II, concerns quickly turned to the use of computers by governments to centralize and concentrate power. These concerns accompanied the expanding use of computers for record keeping and the exponential growth in the scale of databases, allowing the creation, maintenance and manipulation of huge quantities of personal information. This was followed by the inception of software control systems and video games, raising issues of accountability-liability and property rights. This evolution of computer technology can be followed through to more recent developments including the Internet, simulation and imaging technologies, and virtual reality systems. Each one of these developments was accompanied by conceptual and moral uncertainty.

- What will this or that development mean for the lives and values of human beings?
- What will it do to the relationship between government and citizen?
- What will it do to the relationship between employer and employee?
- What will it do to the relationship between businesses and consumers?

According to the sector in which the technology is used

This is to organize the issues according to the sector in which they occur. Ethical issues arise in real-world contexts, and computer-ethical issues arise in the contexts in which computers are used. Each context or sector has distinctive issues and if we ignore this context we can miss important aspects of computer-ethical issues. For example, in dealing with privacy protection in general, we might miss the special importance of privacy protection for *medical records* where confidentiality is so essential to the doctor-patient relationship. Similarly, one might not fully understand the appropriate role for computers in education were one not sensitive to distinctive goals of education. Both of these approaches – examining issues by types and uses of particular technologies, and sector by sector – are important and illuminating; however, they take us too far a field of the philosophical issues.

According to ethical concepts or themes

The third approach is to emphasize ethical concepts and themes that persist across types of technology and sectors. In this the issues are divided into two broad categories:

- Meta-theoretical and methodological issues

- Traditional and emerging issues

Meta-theoretical and methodological issues

Perhaps the deepest philosophical thinking on computer-ethical issues has been reflection on the field itself -- its appropriate subject matter, its relationship to other fields, and its methodology. In a seminal piece entitled "What is Computer Ethics?" Moor (1985) recognized that when computers are first introduced into an environment, they make it possible for human beings (individuals and institutions) to do things they couldn't do before and this creates *policy vacuums*. We do not have rules, policies, and conventions on how to behave with regard to the new possibilities.

- Should employers monitor employees to the extent possible with computer software? Should doctors perform surgery remotely?
- Should I make copies of proprietary software?
- Is there any harm in me taking on a pseudo-identity in an on-line chat room?
- Should companies doing business on-line be allowed to sell the transaction generated-information they collect?

These are examples of policy vacuums created by computer technology. Moor's account of computer ethics has shaped the field of computer ethics with many computer ethicists understanding their task to be that of helping to fill policy vacuums. Indeed, one of the topics of interest in computer ethics is to understand this activity of filling policy vacuums.

2. Philosophical Ethics

Ethics, in general, is categorized in two broad areas – philosophical and professional ethics. We start ethical analysis from philosophical one. To start with, we take unexamined beliefs. But it is important to understand that we can come up with any conclusion after the analysis.

2.1. Ethical Analysis

In philosophical ethical analysis, the reasons for moral beliefs are articulated, and then critically evaluated. The reasons you give for holding an ethical belief or taking a position on an ethical issue can be thought of as an argument for a claim. Once stated, we can ascertain whether the argument does, indeed, support the claim being made or the position being taken. This critical evaluation is often done to reject a position, or to adopt another position, but it may also be done simply to explore a claim. When you critically evaluate the argument supporting a claim, you come to understand the claim more fully. A critical examination of the underpinnings of moral beliefs sometimes leads to a change in belief, but it may also simply lead to stronger and better understood beliefs.

Not only must you give reasons for your claims, you are also expected to be consistent from one argument or topic to the next. For example, instead of having separate, isolated views on abortion and capital punishment, philosophical analysis would lead you to recognize that both your views on abortion and your views on capital punishment rest on a claim about the value of human life and what abrogates it. Philosophical analysis would lead you to inquire whether the claim you made about the value of human life in the context of a discussion of capital punishment is consistent with the claim you made about the value of human life in the context of a discussion of abortion. If the claims appeared to be inconsistent from the one context to the next, then you would be expected to change one of your claims or provide an account of how the two positions can be understood as consistent. In other words, you would show that seemingly inconsistent views are in fact consistent.

2.1.1. Philosophical analysis is a dialectic ongoing process

It involves a variety of activities. It involves expressing a claim and putting forward an argument or reasons for the claim, and it involves critical examination of the argument. If the argument does not hold up to critical examination, then it might be reformulated into a revised argument, perhaps rejecting aspects of the original argument but holding on to a core idea. The revised argument, then, has to be critically examined, and so on, with ongoing reformulation and critique. Philosophers often refer to this process as a dialectic (which is related to the word dialogue). We pursue an argument to see where it goes and to find out what you would have to know or assert to defend the argument and establish it on a firm footing.

In addition to moving from claims to reasons and arguments, and from one formulation of an argument to another better formulation, the dialectic also moves back and forth from cases to principles or theory. To illustrate, take the issue of euthanasia. Suppose you start out by making the claim that euthanasia is wrong. You articulate a principle as the reason for this claim. Say, the principle is that human life has the highest value and, therefore, human life should never be intentionally ended. You might then test this principle by seeing how it applies in a variety of euthanasia cases. For example, is it wrong to use euthanasia when the person is conscious but in extreme pain? When the person is unconscious and severely brain damaged? When the person is terminally ill? When the person is young or elderly? Since your principle concerns the value of human life has implications beyond the issue of euthanasia. Hence, you might also test it by applying it to completely different types of cases. Is intentional taking of human life wrong when it is done in a war situation? Is intentional killing wrong when it comes to capital punishment? Given your position on these cases, you may want to qualify the principle or you may hold to the principle and change your mind about the cases. For example, after seeing how the principle applies in various cases, you may want to qualify it so that you now assert that one should never intentionally take a human life except in self-defence or except when taking a life will save another life. Or you might reformulate the principle so that it specifies that the value of human life has to do with its quality. When the quality of life is significantly and permanently diminished, while it is still not permissible to intentionally kill, it is morally permissible to let a person die.

The dialogue continues as the dialectic leads to a more and more precise specification of the principle and the argument. The process clarifies what is at issue and what the possible positions are. It moves from somewhat inchoate ideas to better and better arguments, and more defensible and better articulated positions.

The dialectic (from an initial belief to an argument, from argument to better argument, and from theory to case, and back) does not always lead to definitive conclusions or unanimous agreement. Therefore, it is important to emphasize that understanding can be improved; progress can be made, even when one has not reached definitive conclusions. Through the dialectic we learn which arguments are weaker and stronger and why. We come to understand the ideas that underpin our moral beliefs. We develop deeper and more consistent beliefs and we come to understand how moral ideas are interrelated and interdependent.

A familiarity with traditional ethical theories will help in articulating the reasons for many of your moral beliefs. Ethical theories provide frameworks in which arguments can be cast. Moreover, ethical theories provide some common ground for discussion. They establish a common vocabulary and frameworks within which, or against which, ideas can be articulated.

2.1.2. Descriptive and Normative Claims

In a sense and partly, this is the distinction between facts and values, but the matter of what counts as a fact is very contentious in philosophy. So, it will be better to stay with the terms descriptive and normative. Descriptive statements are statements that describe a state of affairs in the world. For example, “The car is in the driveway.” And “Georgia is south of Tennessee.” In addressing ethical issues and especially the ethical issues surrounding computer and information technology, it is quite common to hear seemingly factual statements about human beings. The following are descriptive

statements: “Such and such percentage of the people surveyed admitted to having made at least one illegal copy of computer software.” “The majority of individuals who access pornographic Web sites are males between the ages of 14 and 35.” “Such and such percentages of U.S. citizens use the Internet to obtain information on political candidates.” “In all human societies, there are some areas of life that are considered private.” These statements describe what human beings think and do. They are empirical claims in the sense that they are statements that can be verified or proven false by examining the state of affairs described. To be sure, it may not be easy to verify or disconfirm claims like these, but in principle it is possible. Observations can be made, surveys can be administered, and people can be asked, and so on.

Social scientists gather empirical data and report their findings, both on moral and non moral matters. When it comes to morality, psychologists and sociologists might do such things as identify the processes by which children develop moral concepts and sensibilities. Or they may measure how individuals value and prioritize various goods such as friendship, privacy, and autonomy. When anthropologists go to other cultures, they may describe complex moral rules in that culture. They are describing lived and observed moral systems. Similarly, historians may trace the development of a particular moral notion in an historical period.

All of these social scientific studies are descriptive studies of morality; they examine morality as an empirical phenomenon. They don’t, however, tell us what is right and wrong. They don’t tell us what people should think or do, only what people, in fact, think and do.

In contrast, philosophical ethics is normative. The task of philosophical ethics is to explore what human beings ought to do, or more accurately, to evaluate the arguments, reasons, and theories that are proffered to justify accounts of morality. Ethical theories are prescriptive. They try to provide an account of why certain types of behaviour are good or bad, right or wrong. Descriptive statements may come into play in the dialectic about philosophical ethics, but normative issues cannot be resolved just by pointing to the facts about what people do or say or believe. For example, the fact (if it were true) that many individuals viewed copying proprietary software as morally acceptable would not make it so. The fact that individuals hold such a belief is not an argument for the claim that it is morally permissible to copy proprietary software. You might wish to explore why individuals believe this to see if they have good reasons for the belief. Or you might wish to find out what experiences have led individuals to draw this conclusion. Still, in the end, empirical facts are not alone sufficient to justify normative claims. Figuring out what is right and wrong, what is good and what is bad, involves more than a descriptive account of states of affairs in the world.

The aim of this book is not to describe how people behave when they use computers. For this, the reader should consult social scientists—sociologists, anthropologists, political scientists, and psychologists. Rather the aim of this book is to help you understand how people ought to behave when they use computers and what rules or policies ought to be adopted with regard to computer and information technology.

2.2. Ethical Relativism

We can begin our examination of ethical concepts and theories by examining a prevalent, often unexamined moral belief. Many believe that “ethics is relative.” This seems like a good starting place. This claim can be examined carefully and critically. We can begin by formulating the idea as a theory consisting of a set of claims backed by reasons.

The idea of ethical relativism seems to be something like this: “What is right for you may not be right for me,” or “I can decide what is right for me, but you have to decide for yourself.” When we take this idea and formulate it into a more systematic account, it seems to encompass a negative claim (something that it denies), and a positive claim, (something it asserts). The negative claim appears to be: “There are no universal moral norms.” According to this claim, there isn’t a single standard for all human beings. One person may decide that it is right for him to tell a lie in certain

circumstances, another person may decide that it is wrong for her to tell a lie in exactly the same circumstances, and both people could be right. So, the claim that "right and wrong are relative" means in part that there are no universal rights and wrongs.

The positive claim of ethical relativism is more difficult to formulate. Sometimes ethical relativists seem to be asserting that right and wrong are relative to the individual, and sometimes they seem to assert that right and wrong are relative to the society in which one lives. I am going to focus on the latter version, and on this version the relativist claims that what is morally right for me, an American living in the twenty-first century, could be different than what is right for a person living, say, in Asia in the fifth century. The positive claim of relativism is that right and wrong are relative to your society.

Ethical relativists often cite a number of descriptive facts to support these claims:

1. They point to the fact that cultures vary a good deal in what they consider to be right and wrong. For example, in some societies, infanticide is acceptable while in other societies it is considered wrong. In some societies, it is considered wrong for women to go out in public without their faces being covered. Polygamy is permissible in some cultures; in others it is not. Examples of this kind abound.
2. Relativists also point to the fact that the moral norms of a given society change over time so that what was considered wrong at one time, in a given society, may be considered right at another time. Slavery in America is a good example of this since slavery was considered morally permissible by many in the United States at one time, but is now illegal and almost universally considered impermissible.
3. Relativists also point to what we know about how people develop their moral ideas. We are taught the difference between right and wrong as children, and what we come to believe is right or wrong is the result of our upbringing. It depends on when, where, how, and by whom we were raised. If I had been born in certain Middle Eastern countries, I might believe that it is wrong for a woman to appear in public without her face covered. Yet because I was raised in the United States in the twentieth century, by parents who had Western ideas about gender roles and public behaviour, I do not believe this, Of course, parents are not the only determinant of morality. A person develops moral ideas from the experiences he or she has in school, at work, with peers, and so on.

It is useful to note that we have already made progress simply by clearly and systematically formulating the idea of ethical relativism, an idea you may have entertained or heard expressed, but never had a chance to examine carefully. Moreover, we have been able to identify and articulate some reasons thought to support ethical relativism. With the idea and supporting evidence now "on the table," we can carefully and critically examine them.

The facts which ethical relativists point to cannot be denied. For example, I would not want to take issue with the claims that:

1. There is and always has been a good deal of diversity of belief about right and wrong.
2. Moral beliefs change over time within a given society.
3. Social environment plays an important role in shaping the moral ideas you have.

However, there does seem to be a problem with the connection between these facts and the claims of ethical relativism. Do these facts show that there are no universal moral rights or wrongs? Do they show that right and wrong are relative to your society?

On more careful examination, it appears that the facts cited by ethical relativists do not support their claims. To put this another way, we can, without contradiction, accept the facts and still deny ethical relativism. The facts do not necessitate that there are no universal moral standards or that ethics is relative. Lest there be no confusion, you should recognize that "ethics is relative" could be

interpreted either as an empirical or a normative claim. As an empirical claim, it asserts that ethical beliefs vary; as a normative claim it asserts that right and wrong (not just beliefs about, but what is actually right and wrong) vary.

If we understand the claim "ethics is relative" to be a description of human behaviour, then it does follow from the facts sited. Indeed, it is redundant of the facts cited, for as a description of human behaviour, it merely repeats what the facts have said. Ethical beliefs vary. Individuals believe different things are right and wrong depending on how and by whom they have been raised and where and when they live.

On the other hand, if we understand "ethics is relative" to be a normative claim, a claim asserting the negative and/or positive parts of ethical relativism, then it is not redundant, and the facts do not support the claims. Here the leap from facts to conclusion is problematic for a number of reasons. For one, the argument goes from a set of "is" claims to an "ought" claim and the "ought" claim just doesn't follow (in a straightforward way) from the "is" claims.

2.3. Utilitarianism

Utilitarianism is an ethical theory claiming that what makes behaviour right or wrong depends wholly on the consequences. In putting the emphasis on consequences, utilitarianism affirms that what is important about human behaviour is the outcome or results of the behaviour and not the intention a person has when he or she acts. On one version of utilitarianism, what is all important is happiness-producing consequences (Becker and Becker, 1992). Crudely put, actions are good when they produce happiness and bad when they produce the opposite, unhappiness. The term utilitarianism derives from the word utility. According to utilitarianism actions, rules, or policies are good because of their usefulness (their utility) in bringing about happiness.

Lest there be no confusion, philosophers are not always consistent in the way they use the terms utilitarianism and consequentialism. Sometimes, consequentialism is seen as the broadest term referring to ethical theories that claim that what makes an action right or wrong is the consequences and not the internal character of action. Utilitarianism is, then, a particular version of this type of theory with the emphasis specifically on happiness-producing consequences. That is the way I shall use these terms, though I warn readers that the distinction sometimes is made in just the opposite way, that is, with utilitarianism seen as the broadest theory and consequentialism as a particular form of utilitarianism.

In any case, in the version on which I will focus, the claim is that in order to determine what they should do, individuals should follow a basic principle. The basic principle is this: Everyone ought to act so as to bring about the greatest amount of happiness for the greatest number of people.

But, what, you may ask, is the "proof" of this theory? Why should each of us act to bring about the greatest amount of happiness? Why shouldn't we each seek our own interest?

2.3.1. Intrinsic and Instrumental Value

Utilitarians begin by focusing on values and asking what is so important, so valuable to human beings, that we could use it to ground an ethical theory. They note that among all the things in the world that are valued, we can distinguish things that are valued because they lead to something else from things that are valued for their own sake. The former are called instrumental goods and the latter intrinsic goods. Money is a classic example of something that is instrumentally good. It is not valuable for its own sake, but rather has value as a means for acquiring other things. On the other hand, intrinsic goods are not valued because they are a means to something else. They have qualities or characteristics that are valuable in themselves. Knowledge is sometimes said to be intrinsically valuable. So, is art because of its beauty. You might also think about environmental debates in which the value of nature or animal or plant species or ecosystems are said to be valuable independent of

their value to human beings. The claim is that these things have value independent of their utility to human beings.

Having drawn this distinction between instrumental and intrinsic goods, utilitarians ask what is so valuable that it could ground a theory of right and wrong? It has to be something intrinsically valuable, for something which is instrumentally valuable is dependent for its goodness on whether it leads to another good. If you want x because it is a means to y, then y is what is truly valuable and x has only secondary or derivative value. Utilitarianism, as I am using the term, claims that happiness is the ultimate intrinsic good, because it is valuable for its own sake. Happiness cannot be understood as simply a means to something else. Indeed, some utilitarians claim that everything else is de sired as a means to happiness and that, as a result, everything else has only secondary or derivative (instrumental) value.

To see this, take any activity that people engage in and ask why they do it. Each time you will find that the sequence of questions ends with happiness. Take, for example, your career choice. Suppose that you have chosen to study computer science so as to become a computer professional. Why do you want to be a computer professional? Perhaps you believe that you have a talent for computing, and you believe you will be able to get a well-paying job in computer science one in which you can be creative and somewhat autonomous. Then we must ask, why are these things important to you? That is, why is it important to you to have a career doing something for which you have a talent? Why do you care about being well paid? Why do you desire a job in which you can be ca and autonomous? Suppose that you reply by saying that being well paid is important to you because you want security or because you like to buy things or because there are people who are financially dependent on you. In turn, we can ask about each of these. Why is it important to be secure? Why do you want security or material possessions? Why do you want to support your dependents? The questions will continue until you point to something that is valuable in itself and not for the sake of something else. It seems that the questions can only stop when you say you want whatever it is because you believe it will make you happy. The questioning stops here because it doesn't seem to make sense to ask why someone wants to be happy.

A discussion of this kind could go off in the direction of questioning whether your belief is right. Will a career as a computer professional make you happy? Will it really bring security? Will security or material possessions, in fact, make you happy? Such discussions centre on whether or not you have chosen the correct means to your happiness. However, the point that utilitarians

2.3.2. Acts versus Rules

As mentioned earlier, there are several formulations of utilitarianism and proponents of various versions disagree on important details. One important and controversial issue of interpretation has to do with whether the focus should be on rules of behaviour or individual acts. Utilitarians have recognized that it would be counter to overall happiness if each one of us had to calculate at every moment what all the consequences of every one of our actions would be. Not only is this impractical, because it is time consuming and because sometimes we must act quickly, but often the consequences are impossible to foresee. Thus, there is a need for general rules to guide our actions in ordinary situations.

Rule-utilitarians argue that we ought to adopt rules that, if followed by everyone, would, in the long run, maximize happiness. Take, for example, telling the truth. If individuals regularly told lies, it would be very disruptive. You would never know when to believe what you were told. In the long run, a rule obligating people to tell the truth has enormous beneficial consequences. Thus, "tell the truth" becomes a utilitarian moral rule. "Keep your promises," and "Don't reward behaviour that causes pain to others," are also rules that can be justified on utilitarian grounds. According to rule-utilitarianism, if the rule can be justified in terms of the consequences that are brought about from people following it, then individuals ought to follow the rule.

Act-utilitarians put the emphasis on individual actions rather than rules. They believe that even though it may be difficult for us to anticipate the consequences of our actions, that is what we should be trying to do. Take, for example, a case where lying may bring about more happiness than telling the truth. Say you are told by a doctor that tentative test results indicate that your spouse may be terminally ill. You know your spouse well enough to know that this knowledge, at this time, will cause your spouse enormous stress. He or she is already under a good deal of stress because of pressures at work and because someone else in the family is very ill. To tell your spouse the truth about the test results will cause more stress and anxiety, and this stress and anxiety may turn out to be unnecessarily if further tests prove that the spouse is not terminally ill. Your spouse asks you what you and the doctor talked about. Should you lie or tell the truth? An Act-utilitarian might say that the right thing to do in such a situation is to lie, for little good would come from telling the truth and a good deal of suffering (per haps unnecessary suffering) will be avoided from lying. A rule-utilitarian would agree that good might result from lying in this one case, but in the long run, if we cannot count on people telling the truth (especially our spouses), more bad than good will come. Think of the anxiety that might arise if spouses routinely

2.4. Deontological Theories

In utilitarianism, what makes an action or a rule right or wrong is outside the action; it is the consequences of the action or rule that make it right or wrong. By contrast, deontological theories put the emphasis on the internal character 'of the act itself.' What makes an action right or wrong for deontologists is the principle inherent in the action. If an action is done from a sense of duty, if the principle of the action can be universalized, then the action is right. For ex ample, if I tell the truth (not just because it is convenient for me to do so, but) because I recognize that I must respect the other person, then I act from duty and my action is right. If! tell the truth because I fear getting caught or because I believe I will be rewarded for doing so, then my act is not morally worthy.

I am going to focus here on the theory of Immanuel Kant. If we go back for a moment to the allocation of dialysis machines, Kant's moral theory is applicable -because it proposes what is called a categorical imperative specifying that we should never treat human beings merely as means to an end. We should always treat human beings as ends in themselves. Although Kant is not the only deontologist, I will continue to refer to him as I discuss deontology.

The difference between deontological theories and consequentialist theories was illustrated in the discussion of allocation of dialysis machines. Deontologists say that individuals are valuable in themselves, not because of their social value. Utilitarianism is criticized because it appears to tolerate sacrificing some people for the sake of others. In utilitarianism, right and wrong are dependent on the consequences and therefore vary with the circumstances. By contrast, deontological theories assert that there are some actions that are al ways wrong, no matter what the consequences. A good example of this is killing. Even though we can imagine situations in which intentionally killing one person may save the lives of many others, deontologists insist that intentional killing is always wrong. Killing is wrong even in extreme situations be cause it means using the person merely as a means and does not treat the human being as valuable in and of himself. Deontologists do often recognize self-defence and other specia l circumstances as excusing killing, but these are cases when, it is argued, the killing is not exactly intentional. (The person attacks me. I would not, otherwise, aim at harm to the person, but I have no other choice but to defend myself.)

At the heart of deontological theory is an idea about what it means to be a person, and this is connected to the idea of moral agency. Charles Fried (1978) put the point as follows:

Substantive contents of the norms of right and wrong express the value of persons, of respect for personality. What we may not do to each other, the things which are wrong, are precisely those forms of personal interaction which deny to our victim the status of a freely choosing, rationally valuing, specially efficacious person, the special status of moral personality.

According to deontologists, the utilitarians go wrong when they fix on happiness as the highest good. Deontologists point out that happiness cannot be the highest good for humans. The fact that we are rational beings, capable of reasoning about what we want to do and then deciding and acting, suggests that our end (our highest good) is something other than happiness. Humans differ from all other things in the world insofar as we have the capacity for rationality. The behaviour of other things is determined simply by laws of nature. Plants turn toward the sun because of photosynthesis. They don't think and decide which way they will turn. Physical objects fall by the law of gravity. Water boils when it reaches a certain temperature. In contrast, human beings are not entirely determined by laws of nature. We have the capacity to legislate for ourselves. We decide how we will behave. As Kant describes this, it is the difference between acting in accordance with law (plants and stones do) and acting in accordance with the conception of law.

The capacity for rational decision making is the most important feature of human beings. Each of us has this capacity; each of us can make choices, choices about what we will do, and what kind of persons we will become. No one else can or should make these choices for us. Moreover, we should recognize this capacity in others.

Notice that it makes good sense that our rationality is connected with morality, for we could not be moral beings at all unless we had this rational capacity. We do not think of plants or fish or dogs and cats as moral beings precisely because they do not have the capacity t9 reason about their actions. We are moral beings because we are rational beings, that is, because we have the capacity to give ourselves rules (laws) and follow them.

Where utilitarians note that all humans seek happiness, deontologists emphasize that humans are creatures with goals who engage in activities directed toward achieving these goals (ends), and that they use their rationality to formulate their goals and figure out what kind of life to live. In a sense, deontologists pull back from fixing on any particular value as structuring morality and instead ground morality in the capacity of each individual to organize his or her own life, make choices, and engage in activities to realize their self-chosen life plans. What morality requires is that we respect each of these beings as valuable in themselves and refrain from valuing them only insofar as they fit into our own life plans.

As mentioned before, Kant put forward what he called the categorical im /Jerutive. While there are several versions of it, I will focus on the second version, which goes as follows: Never treat another human being merely as a means but always as an end. This general rule is derived from the idea that persons are moral beings because they are rational, efficacious beings. Because we each have the Capacity to think and decide and act for ourselves, we should each be treated with respect, that is, with recognition of this capacity.

Note the "merely" in the categorical imperative. Deontologists do not in Sift that we never use another person as a means to an end, only that we never merely use them in this way. For example, if I own a company and hire employees to work in my company, I might be thought of as using those employees as a means to my end (i.e., the success of my business). This, however, is not wrong if I promise to pay a fair wage in exchange for work and the employees agree to work for me. I thereby respect their ability to choose for themselves.

2.5. Rights

So far, very little has been said about rights though we often use the language of rights when discussing moral issues. "You have no right to tell me what to do." "I have a right to do that." Ethicists often associate rights with deontological theories. The categorical imperative requires that each person be treated as an end in himself or herself, and it is possible to express this idea by saying that individuals have "a right to" the kind of treatment that is implied in being treated as an end. The idea that each individual must be respected as valuable in himself or herself implies that we each

have rights not to be interfered with in certain ways, for example, not to be kille d or enslaved, to be given freedom to make decisions about our own lives, and so on.

An important distinction that philosophers often make here is between negative rights and positive rights. Negative rights are rights that require restraint by others. For example, my right not to be killed requires that others refrain from killing me. It does not, however, require that others take positive action to keep me alive. Positive rights, on the other hand, imply that others have a duty to do something to or for the right holder. So, if we say that I have a positive right to life, this implies not just that others must refrain from killing me, but that they must do such things as feed me if! am starving, give me medical treatment if! am sick, swim out and save me if! am drowning, and so on. As you can see, the difference between negative and positive rights is quite significant.

Positive rights are more controversial than negative rights because they have implications that are counter-intuitive. If every person has a positive right to life, this seems to imply that each and every one of us has a duty to do what ever is necessary to keep all people alive. This would seem to suggest that, among other things, it is our duty to give away any excess wealth that we have to feed and care for those who are starving or suffering from malnutrition. It also seems to imply that we have a duty to supply extraordinary life-saving treatment for all those who are dying. In response to these implications, some philosophers have argued that individuals have only negative rights.

While, as I said earlier, rights are often associated with deontological theories, it is important to note that rights can be derived from other theories as well. For example, we can argue for the recognition of a right to property on utilitarian grounds. Suppose we ask why individuals should be allowed to have private property in general and, in particular, why they should be allowed to own computer software. Utilitarians would argue for private ownership of soft ware on grounds that much more and better software will be created if individuals are allowed to own (and then license or sell) it. Thus, they argue that individuals should have a legal right to ownership in software because of the beneficial consequences of acknowledging such a right.

Another important thing to remember about rights is the distinction between legal and moral (or natural or human) rights. Legal rights are rights that are created by law. Moral, natural, or human rights are claims independe nt of law. Such claims are usually embedded in a moral theory or a theory of human nature.

The utilitarian argument is an argument for creating or recognizing a legal right; it is not an argument to the effect that human beings have a natural right, for example, to own what they create.

2.5.1. Rights and Social Contract Theories

Rights are deeply rooted in the tradition of social contract theories. In this tradition the idea of a social contract (between individuals, or between individuals and government) is hypothesized to explain and justify the obligations that human beings have to one another. Many of these theories imagine human beings in a state of nature and then show that reason would lead individuals in such a state to agree to live according to certain rules, or to give power to a government to enforce certain rules. The depiction of a state of nature in which human beings are in a state of insecurity and uncertainty is used to suggest what human nature is like and to show that human nature necessitates government. That is, in such a state any rational human beings would agree (make a contract) to join forces with others even though this involves giving up some of their natural freedom. The agreement (the social contract) creates obligations and these are the basis of moral obligation.

An argument of this kind is made by several social contract theorists and each specifies the nature and limits of our obligations differently. One important difference, for example, is in whether morality exists prior to the social contract. Hobbes argues that there is no justice or injustice in a state of nature; humans are at war with one another and each individual must do what they must to preserve themselves. Locke, on the other hand, specifies a natural form of justice in the state of nature. Human beings have rights in the state of nature and others can treat individuals unjustly.

Government is necessary to insure that natural justice is implemented properly because without government, there s no certainty that punishments will be distributed justly.

2.5.2. Rawlsian Justice

In 1971,John Rawls, a professor at Harvard University, introduced a new version of social contract theory (though some argue it is not a social contract theory in the traditional sense). Rawls introduced the theory in a book entitled simply A Theory of Justice. The theory may well be one of the most influential theories of the twentieth century, for not only did it generate an enormous amount of attention in the philosophical community, it influenced discussion among economists, social scientists, and public policy makers.

Rawls was primarily interested in questions of distributive justice. In the tradition of a social contract theorist, he tries to understand what sort of contract between individuals would be just. Rawls recognizes that we can't arrive at an account of justice and the fairness of social arrangements by reasoning about what rules particular individuals would agree to. He understands that individuals are self-interested and therefore will be influenced by their own experiences and their own situation when they think about fair arrangements. Thus, if some group of us were to get together in something like a state of nature (suppose a group is stranded on an island or a nuclear war occurs and only a few survive), the rules we would agree to would not necessarily be a just system. It would not necessarily exemplify justice.

The problem is that we would each want rules that would favour us. Smart people would want rules that favoured intelligence. Strong people would want a System that rewarded strength. Women would not want rules that were biased against women, and so on. The point is that there is no reason to believe that the outcome of a negotiation in which people expressed their preferences would result in rules of justice and just institutions. In this sense, Rawls believes that justice has to be blind in a certain way.

Rawls specifies, therefore, that in order to get at justice, we have to imagine that the individuals who get together to decide on the rules for society are behind a veil of ignorance. The veil of ignorance is such that individuals do not know what characteristics they will have. They do not know whether they will be male or female, black or white, highly intelligent or moderately intelligent or retarded, physically strong or in ill-health, musically talented, successful at business, indigent and so on.

At the same time, these individuals would be rational and self-interested and would know something about human nature and human psychology. In a sense, what Rawls is suggesting here is that we have to imagine generic human beings. They have abstract features that human beings generally have (i.e., they are rational and self-interested). And, they have background knowledge (i.e., general knowledge of how humans behave and interact and bow they are affected in various ways).

According to Rawls, justice is what individuals would choose in such a situation. Notice that what he has done, in a certain sense, is eliminate bias in the original position. Once a society gets started, once particular individuals have characteristics, their views on what is fair are tainted. They cannot be objective.

So, justice, according to Rawls is what people would choose in the original position where they are rational and self-interested, informed about human nature and psychology but behind a veil of ignorance with regard to their own characteristics. Rawls argues that individuals in the original position would agree to two rules. These are the rules of justic e and they are "rules of rules" in the sense that they are general principles constraining the formulation of specific rules. The rules of justice are:

1. Each person should have an equal right to the most extensive basic liberty compatible with a similar liberty for others.

2. Social and economic inequalities should be arranged so that they are both (a) reasonably expected to be to everyone's advantage and (b) attached to positions and offices open to all.

These general principles assure that no matter where an individual ends up in the lottery of life (in which characteristics of intelligence, talents, physical abilities, and so on, are distributed), he or she would have liberty and opportunity. He or she would have a fair shot at a decent life.

While Rawls' account of justice has met with criticism, it goes a long way in providing a framework for envisioning and critiquing just institutions. This discussion of Rawls is extremely abbreviated as were the accounts of Kant and utilitarianism. Perhaps the most important thing to keep in mind as we

2.6. Virtue Ethics

Before moving on to the ethical issues surrounding computer and information technology, one other tradition in ethical theory should be mentioned. In recent ears, interest has arisen in resurrecting the tradition of virtue ethics, a tradition going all the way back to Plato and Aristotle. These ancient Greek philosophers pursued the question: What is a good person? What are the virtues associated with being a good person? For the Greeks virtue meant excellence, and ethics was concerned with excellences of human character. A person possessing such qualities exhibited the excellences of human good. To have these qualities is to function well as a human being.

The list of possible virtues is long and there is no general agreement on which are most important, but the possibilities include courage, benevolence, generosity, honesty, tolerance, and self-control. Virtue theorists try to identify e list of virtues and to give an account of each - what is courage? What is honesty? They also give an account of why the virtues are important.

Virtue theory seems to fill a gap left by other theories we considered, because it addresses the question of moral character, while the other theories focused primarily on action and decision making. What sort of character should we be trying to develop in ourselves and in our children? We look to moral heroes, for example, as exemplars of moral virtue. Why do we admire such people? What is it about their character and their motivation that are worthy of tour admiration?

Virtue theory might be brought into the discussion of computer technology and ethics at any number of points. The most obvious is, perhaps, the discussion of professional ethics, where we want to think about the characteristics of a good computer professional. Good computer professionals will, perhaps, exhibit honesty in dealing with clients and the public. They should exhibit courage when aced with situations in which they are being pressured to do something illegal or act counter to public safety. A virtue approach would focus on these characteristics and more, emphasizing the virtues of a good computer professional.

2.7. Individual and Social Policy Ethics

One final distinction will be helpful. In examining problems or issues, it is important to distinguish levels of analysis, in particular that between macro and micro level issues or approaches. One can approach a problem from the point of view of social practices and public policy, or from the point of view of individual choice. Macro level problems are problems that arise for groups of people, a community, a state, a country. At this level of analysis, what is sought is a solution in the form of a law or policy that specifies how people in that group or society ought to behave, what the rules of that group ought to be. When we ask the following questions, we are asking macro level questions: Should the United States grant software creators a legal right to own software? Should software engineers be held liable for errors in the software they design? Should companies be allowed to electronically monitor their employees?

On the other hand, micro level questions focus on individuals (in the presence or absence of law or policy). Should I make a copy of this piece of soft ware? Should I lie to my friend? Should I work on

a project making military weapons? Sometimes these types of questions can be answered simply by refer ring to a rule established at the macro level. For example, legally I can make a back-up copy of software that I buy, but I shouldn't make a copy and give it to my friend. Other times, there may be no macro level rule or the macro level rule may be vague or an individual may think the macro level rule is unfair. In these cases, individuals must make decisions for themselves about what they ought to do.

The theories just discussed inform both approaches, but in somewhat different ways, so it is important to be clear on which type of question you are asking or answering.

3. Social Context of a Design

We only have identified the primary factors in a computer system to be performance, reliability, and cost. The factors such as usability and fit-to-task figuring in any software system are not considered by us. A limited effort is included in the analysis of social context in tra ditional software design. One of the central challenges faced by software designers is how to balance the highly structured nature of computer artifacts with the need to integrate them into different settings. A software designer inevitably faces situations in which design choices are constrained by:

- The conflicting goals and values held by the different parties who have a stake in the changes that new technologies will bring to the work
- Further, Workers and managers have many
    - Common interests
    - Different stakes

The question will be how computers in the workplace change productivity, working conditions, and job satisfaction.

3.1. Design for People at Work

Well-designed systems can boost

- Productivity
- Enhance
- Job satisfaction and give both workers and managers a clearer sense of what is going on in the organization.

But a system that interferes with crucial work practices:

- Can result in reduced effectiveness and efficiency,
- Reduced satisfaction and autonomy,
- Increased stress and health problems for the people who use the system.

3.2. Design Approaches

The basic design approaches are

- Technology-centered approach
- Work-Oriented Approaches

Work-Oriented approaches are further classified into two design approaches:

- Human-centered design

- Participatory design

3.2.1. Technology-centered approach

This approach is characterized by hard-systems thinking; which involve the imposition of a clear-cut problem definition on a relatively unstable organizational reality and a "fuzzy" system.

In this design a little attention is accorded to the organizational context in which the system is to operate or the social implications of the system.

Human-centered design

This design puts human, social, and organizational considerations on an equal footing with technical considerations in the design process. Well-designed technology should make use of human strengths such as skill, judgment, capacity for learning—to create a robust and flexible production system.

Participatory design

This approach emphasizes the importance of meaningful end-user participation and influence in all phases of the design process. "The employees must have access to relevant information; they must have the possibility for taking an independent position on the problems; and they must in some way participate in the process of decision making."

3.3. Systems and Assumptions

Consider the following three cases with a view of assumptions made when bringing these systems into use in a particular context. Many assumptions are ruled out because the output is not what is expected. This gives us more understanding in relation to the topic discussed above that fit to task and usability are also important parameters to be considered in designing a system.

Case 1: The Trouble Ticketing System

The Trouble Ticketing System (TTS) was a mainframe-based system, developed in the early 1980s for use by telephone-company repair personnel, for scheduling, work routing, and record keeping. When trouble was reported, a job ticket was generated, and was sent to the appropriate telephone company office. There, a worker picked up the ticket and began work on the job. When the repair was completed, or when the worker had done all that she could do from that location, the ticket was sent back to the central TTS.

Before TTS, job tickets were generated, but work was more collaborative. Testers called one another, consulting with someone at the other end of a problematic line, or someone who knew a particular part of the system especially well, for help with troubleshooting. One of the motivations for the development of the TTS was to ensure that workers spent more time on repair tasks by eliminating conversations between workers, which were thought to be inefficient and "off task." With the TTS, each tester worked alone. If she could not complete a repair job, the tester recorded what had been done and sent the ticket back to the TTS for someone else to pick up and work on.

The need for conversation was eliminated, but the benefits of conversation (more information available to diagnose problems, and the chance to learn more about the system) were lost too. Because TTS also monitored the number of hours that each worker spent doing jobs, testers who spent time consulting with one another or training new workers (neither of which activities were accounted for in the system) were penalized.

While TTS was designed to make job performance more efficient, it has created the opposite effect: discouraging the training of new hands, breaking up the community of practice by eliminating troubleshooting conversations, and extending the time spent on a single job by segmenting coherent troubleshooting efforts into unconnected ticket-based tasks

Case 2: Big Bank

At a large urban bank, teller operations were supported by a computer-based system, which was designed and modified over time by the bank's systems department. The system had to meet the bank's needs for accuracy, efficiency, security, and customer service, although, as this case illustrates, not all of these goals can maximized simultaneously. The Big Bank system embodied in a rigid form the rules and procedures of the bank, which before automation would have been enforced by people--generally supervisors--who were relatively close to the action and who exercised professional judgment based on experience with the day-to-day work of bank tellers. One of the security features built into the system was that, "under a specified set of exceptional situations (defined by management at the bank through the setting of software parameter settings), the teller's terminal will freeze with a message about the account on the screen." The transaction could be completed only after a bank officer's authorization card was passed through a card reader on the teller's terminal. The purpose of the freeze was to ensure supervisory oversight in circumstances that were deemed exceptional, such as more than three transactions on a single account on a single day.

It turned out, however, that at a large downtown branch of the bank, the exceptional happened every 5 or 10 minutes. To keep lines flowing and to avoid costly inefficiencies, the manager gave a bank officer's card to the tellers, which they passed among themselves to unfreeze their terminals. Only once or twice an hour did the tellers judge that a supervisor was needed--when a customer was unknown, or when an unusual situation arose. Of course, overcoming this one feature undermined all security provisions in the system, because the officer card was now freely available to the tellers. "This implementation of the system increases the responsibility of the teller although, at the same time, the design of the system is reminding the teller that, formally, the bank management does not trust him or her to make even routine decisions. The result of the system's design is the worst of both worlds."

Case 3: HELP System

In a large machining area in a production plant, a major U.S. aircraft manufacturer installed a new computer-based system known as the HELP (Help Employees Locate People) system. The system had two principal functions. The first--which was the source of its official name--was to enable machinists to signal for assistance when they needed replacement tooling, wanted consultation, or were due for a break. At the push of a button, a machinist could indicate his location and the nature of the request. This aspect of the system got good reviews from both machinists and shop management, because it enabled operations to run more smoothly.

The second function of the system, not acknowledged in the official name, was the monitoring of 66 machines in the shop. A panel in a control room above the shop floor displayed the status of each machine with colored lights. A supervisor could check these lights, and could gain further information by glancing out at the floor below. Daily reports told supervisors not only about production levels, but also about how each worker spent his time. Upper management received weekly and monthly reports.

The purpose of the monitoring system was to gain greater managerial control over how machinists spent their time, as an aid to increasing productivity. The problem was that the information captured by the system and reported to the managers provided only a partial view of the work of machinists--a view so one-sided as to be nearly useless. The very concreteness of the statistics, however, invited unwarranted conclusions. For example, managers wanted to discourage machinists from slowing down production unnecessarily, so the system was designed to report all the time that a computer-controlled machine tool spent halted or operating at less than 80 percent of the programmed feed rate (the rate at which the cutting tool moves across the surface of the metal). What the system did not report, however, was whether the programmer who wrote the program for that particular part had set the feed rate correctly. It is common for feed rates to be set incorrectly, and one of the skills of the machinist is to judge whether conditions require slowing down or allow speeding up. Information about the appropriate feed rate is crucial to the proper interpretation of the data generated by the monitoring system.

The data generated by the HELP system could mislead managers into thinking that any machine status other than "running" was an indication of unproductive activity. "One machinist...had to work long and furiously to set up a particularly intricate part to be cut. As a result, his machine sat idle most of the day. While he felt that he had never worked harder, his supervisor reprimanded him because the system reported that his machine was idle."

The sense of a hostile, intrusive presence in machinists' work life gave the system its unofficial shop floor name: The Spy in the Sky. The system was tolerated by the machinists because of its benefits as a signaling

device, and because their union was able to negotiate an agreement with the company that data from the system would not be used for disciplinary purposes. Still, the system could be used for informal discipline, and it served as a constant reminder of managers' mistrust of the workforce. Productive work relies on the skills of machinists, yet the monitoring system embodies the assumption that, if machines are running at less than 80 percent of the programmed rate, the fault lies with the machinist, rather than with the engineer who wrote the parts program.

We see form all the above cases that the basic intent of bringing a computer system into an existing scenario does not exactly meet all the assumptions with which it is brought. The implementation basically sees the following impact.

- **Representations and Misrepresentations of Work**
- **The High Cost of Bad Design**
- **Displacement.**
- **Intensification.**
- **Reduction or redefinition of skill**

Thus to conclude the social context of a design should also consider the usability and fit to task scenario when designing or redefining the existing systems with computer systems. We should not attach false assumptions and take more realistic measures into consideration.

Source:

Deborah G. Johnson, *Computer Ethics*, Third Edition